Submission in Response to NSF CI 2030 Request for Information

DATE AND TIME: 2017-04-05 15:49:18
REFERENCE NO: 267

Author Names & Affiliations

- Gregory Peterson - University of Tennessee/National Institute for Computational Sciences

Contact Email Address (for NSF use only)

(Hidden)

Research Domain, discipline, and sub-discipline

Computational and Data Science and Engineering

Title of Submission

Future Cyberinfrastructure Needs to the Computational and Data Science and Engineering Community

Abstract (maximum ~200 words).

The scientific and engineering community continues to pursue revolutionary improvement in our understanding of the universe and in creating exciting new technologies to improve our lives. Computational and data science and engineering helps deliver these promises, but advanced cyberinfrastructure improvement is required to continue making progress.

Question 1 Research Challenge(s) (maximum ~1200 words): Describe current or emerging science or engineering research challenge(s), providing context in terms of recent research activities and standing questions in the field.

The scientific and engineering community continues to pursue revolutionary improvement in our understanding of the universe and in creating exciting new technologies to improve our lives. Computational and data science and engineering helps deliver these promises, but advanced cyberinfrastructure improvement is required to continue making progress.

The following list includes some of the most noteworthy opportunities and challenges facing the national research community.

- Efficient scaling to exascale for capability computing
- Efficient scaling for large ensembles (HTC)
- Scalable data life cycle support
- Reproducibility
- Researcher workforce development
- Applications development workforce development
- CI operations workforce development

Clearly, adequate CI capabilities must be made available to the community for economic development and scientific inquiry to continue. At the lower end of the spectrum, computational resources are ubiquitous for individuals. Delivering exascale and beyond computing systems will be required for capability computing that addresses grand challenge problems and drives future computer system development. The computing systems that are “in between” must be provisioned as well to enable problems to scale up over time so that tomorrow’s grand challenge problems can be envisioned and tackled. Whether traditional modeling and simulation workloads and emerging analytics workloads can effectively target the same architectures or will demand separate systems remains to be seen. Similarly, high throughput computing needs may be largely supported by cloud computing resources, although this is a rapidly changing technology. Supporting this diverse set of computational resources presents tremendous challenges to the community. For example, although the NSF has recently sought to target new communities with HPC resources, traditional modeling and simulation users now face a dearth of platforms for their computational resource.

Radical shifts can occur across computing system scales with respect to programming, usage, administration, and even physical deployment. For programming and application development, this implies not only changes to algorithms that scale better, but potential additional complexities associated with different languages, compilers, debuggers, and optimization tools. Users must also learn to navigate disparate allocations processes to obtain access and then tune their workflows for different queuing systems and policies with respect to computation along with data movement and storage. As key components of the cyberinfrastructure, system administrators face tremendous challenges with deploying and supporting systems across the spectrum of resources, in addition to managing security, privacy, sensitive data, and policy issues.

A coordinated national cyberinfrastructure strategy is required to deliver what the scientific and engineering community needs. The coordination should include support for computational resources and access to them (e.g., with better communications and coordination between allocations processes), a coordinated national strategy for supporting the data life cycle for data products and software, and workforce development that addresses computational researchers, applications developers, and cyberinfrastructure operational staff.

The community requires coordinated development to enable the successful targeting and exploitation of the entire spectrum described in the “Branscomb pyramid” of resources. Of particular concern is coordinated workforce development for researchers, application developers, and operational staff to climb the pyramid. The community needs a coordinated plan for how scientific and engineering researchers can develop their capabilities to effectively target small to grand challenge problems using laptops to leadership class supercomputers. Accordingly, application development work that helps researchers to seamlessly target this spectrum of computational platforms is needed. Key challenges for application development includes the changes to computer architectures and technologies (e.g., support for GPUs and MICs), changes to languages and programming models, and programming techniques that can scale from tens to billions of threads of execution. The explosion of computing now represents a significant national challenge in training a sufficient number of operational staff that can support the spectrum of computing systems. Viable professional development and career pathways for operational staff represents a key challenge to the community and should be addressed in a coordinated way across federal agencies.

Question 2 Cyberinfrastructure Needed to Address the Research Challenge(s) (maximum ~1200 words): Describe any limitations or absence of existing cyberinfrastructure, and/or specific technical advancements in cyberinfrastructure (e.g. advanced computing, data infrastructure, software infrastructure, applications, networking, cybersecurity), that must be addressed to accomplish the identified research challenge(s).

See question 1 response.

Question 3 Other considerations (maximum ~1200 words, optional): Any other relevant aspects, such as organization, process, learning and workforce development, access, and sustainability, that need to be addressed; or any other issues that NSF should consider.

See question 1 response.

Consent Statement

- "I hereby agree to give the National Science Foundation (NSF) the right to use this information for the purposes stated above and to display it on a publically available website, consistent with the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode)."